
PROFORMA GLOBAL RESEARCH

Your Board Wants AI. Where Do You Start?

Matt Rollings, Founder and Principal, Proforma Global

Briefing · 2026-06-06

The board is now pushing for AI rollouts and adoption, and the timer has started. Failed AI rollouts will drastically alter career trajectories.

MIT's 2025 research found that 95% of enterprise AI deployments produce no measurable return, and without a plan, your organization will assuredly be a part of that 95%.

Most guidance is obvious. We try to highlight the non-obvious.

The Obvious

The obvious that you already know:

- "Deliver AI" is not a goal. The board is looking for ROI and a date to measure against.
- You need to hire a leader who has experience deploying AI to make this work.
- Pilots should target use cases with stable processes.
- Build internal capability rather than a standing dependency on a consultant. Work that leaves the dependency in place has transformed nothing.
- Successful AI deployments require a significant commitment to data cleansing.
- Successful AI transformations require a complete reengineering of your processes.

Genuine Insight

The program is won or lost on a handful of decisions that rarely make the roadmap.

Financial Systems require a completely different deployment framework. Don't start there.

The most tempting target for AI Agent deployment is within the financial systems frameworks: ERP, FP&A tools (EPM), etc. The issue is that AI requires a completely different approach that effectively no one has solved yet.

Financial systems require any consumer of information to handle the chart of accounts segments: Account, Company, Cost Center, Product, etc. The term for this is multidimensional. AI frameworks do not process multidimensional data well, and it requires a highly structured set of supporting layers for the systems to read it.

Starting with Financial Systems is tempting due to the ROI, but it likely will be the most complex rollout you have. Start smaller and build competency before tackling this.

Scaling AI Agent frameworks within financial systems has not yet been successfully achieved

Achieving long-term success with AI in the ERP/EPM worlds requires a complete change of process and a different approach to the technology.

The current design thinking:

1. Does not scale
2. Is not auditable
3. Is inherently ROI-negative
4. Is prone to hallucination

Additionally, no available 'off-the-shelf' tools can currently support an enterprise deployment. Each solution must be custom-built based on implementer knowledge, predicated on the skill of the project team, not the brand.

AI agent frameworks within financial systems are 15–20% AI and 80% dynamic automation.

Most AI frameworks are failing at scale because AI usage is 80% of the solution. The allure of a fully dynamic universal tool to support any process is a mirage that cannot work with the current models.

Financial systems data is highly complex and prone to overloading AI with the required relationships and context. The sole predictor of success of AI deployments within financial systems is a mature data architecture. This fact is not true for other areas, due to the lower complexity of their data.

Most self-proclaimed experts are novices

The main issue with AI is that those claiming to be experts are novices in its actual deployment strategies.

Successful AI deployments require:

1. Deep technical understanding.
2. Deep expertise in operations and process design.
3. Subject matter expertise of the systems it is deployed onto.

4. Deep transformation background.
5. Deep implementation expertise across at least one major software package.

Most of the information being produced on how to implement AI is correct academically, but when evaluated against complex data and multidimensional systems, the reality is starkly different. Even the AI experts (Anthropic/OpenAI) are not aware of how to deploy AI in a sustainable/scalable fashion across organizations. Their financial systems are still in their infancy, and deploying a proper scalable foundation takes years to mature.

Only a handful of true experts exist across each domain, with many novices placed into leadership roles trying to guide organizations on their journey. The intent is not to do harm, but we all know how consulting is: fake it until you make it, and collect the revenue while you can.

The career risk for a failed deployment due to an improper strategy has never been higher.

AI Agents are intended to be operations specialists

AI Agents are not ad-hoc tools. They replace discrete tasks within operational workflows currently done by humans. Additionally, agent frameworks are not chatbots. Chatbots to get information from a database may be attractive, but running a financial report from a modern tool is many orders of magnitude more efficient by any metric.

Generally, AI Agents are intended to augment and redesign how each function operates. Tasks within the Office of the CFO are broken into two categories:

1. Highly structured, repeatable processes. These are best solved by agentic AI.
2. Ad-hoc tasks that require changing business assumptions with varying levels of intent (e.g. what-if modeling).

The ad-hoc tasks within a financial systems framework are not conducive to automation through Agents.

An organization can absolutely use an interactive AI model to help create financial models, but that is not an AI Agent. Ad-hoc requests can be completed with the assistance of Claude and Claude Code, but those are not agent deployments. However, an internal resource that is a novice with AI can spend upwards of \$5,000 per day in tokens creating and analyzing financial models.

The uncomfortable truth is that a significant portion of Finance as a function is ad-hoc and is not entirely suitable for AI Agents. A different type of AI deployment is required.

The question here is not whether an AI can perform a task, but whether an operator can work with an AI tool to create a utility to generate output to solve an operational problem reliably, in a cost-efficient manner. Can that utility be handed off to another resource and not create the decentralized Excel nightmare that plagued all companies at one point or another?

AI frameworks should focus on automating and augmenting repeatable, stable processes. That is the opposite use case from what interactive chatbots are for. The more 'chatbot' functionality the agent has, the less reliable it will be.

Hire an independent liaison to help you guide larger implementers

The real key to success is to cut through the noise and accept that sometimes there are not enough educated people internally to know whether the decisions being made are good. Consulting partners are not always going to be honest with you, due to conflicts of interest.

Onboarding an independent resource to help advise your team and make the right decisions is critical. Identifying risks during the selling and contracting phase is arguably the most important element you have. Even with an internal AI architect, that person is oftentimes constrained by politics. Success in this space is entirely predicated on knowledge. Don't risk the outcome purely based on a lack of expertise or internal political maneuvering.

Related reading:

- [Implementing Your First AI Agent Pilot](#) - how to execute the first effort once you have picked a target.
- [Semantic Layers in Enterprise Agent Systems](#) - why multidimensional financial data needs a structured layer before an agent can read it.
- [Data Architecture for Enterprise Agents](#) - the data foundation that decides whether a deployment scales.
- [The Agentic AI Maturity Curve in Enterprise Finance](#) - which architecture fits which stage.

[DOWNLOAD PDF](#)

Work with Proforma Global

The thinking in this paper is public; the methods that turn it into a working system are not. If it fits a problem your team is working on, that is what we bring to an engagement. Start a conversation: info@proforma.global.

© 2026 Proforma Global. All rights reserved.

This paper is published as Proforma Global Research. The text and figures are the property of Proforma Global.

Brief excerpts may be quoted under fair use with attribution to Proforma Global Research and a link to the canonical URL. Permission requests: info@proforma.global.